

Chassis Cluster Upgrade with Minimal Downtime (v1.0)

SRX BRANCH (SRX1xx, SRX2xx, SRX3xx, SRX550, SRX550HM, SRX650)

SRX Mid-Range (SRX1400, SRX1500, SRX3400, SRX3600, SRX4100, SRX4200, SX4600, vSRX)

Prerequisites:

- Prerequisites may be performed outside of MW with no impact to traffic
- Console connections setup to both chassis cluster nodes are necessary to allow unique config adjustments and due to device power off via 'halt' method used
- Download Junos software from [Juniper Download website](#)
- Backup current configuration
 - Local device file storage – 'show configuration | save /var/tmp/<name>
 - To attached USB – Use [KB12880](#) to mount USB then 'save configuration | save /var/tmp/usb'
- Upload Junos OS image to Device storage
 - e.g., /var/tmp/junos-install-srx5000-x86-64-18.4R2.7.tgz
- Verify upgrade image package to current configuration
 - Not available for SRX1500, SRX4100/4200, SRX4600 and vSRX devices
 - >request system software validate <image location>
- Temporarily disable connected switch settings for MAC moves/duplications such as 'mac-move-limit' and 'duplicate-mac-detection' due to possible duplicate location mac addresses during Step 12.

Upgrade Directions:

The below steps assume that Node0 is the primary for control plane (RG0) and data plane (RG1+) and configured with a higher priority than the secondary node.

As needed, please failover all redundancy-groups (RGs) to primary node
>request chassis cluster failover redundancy-group [x] node 0

Node0 Directions	Node1 Directions
<p>1. Disable all physical interfaces used for transit traffic on Node1 (secondary node) <i>Note: Alternatively, you may physically remove cables or 'shut' connected device interfaces.</i></p> <p>e.g.,</p> <pre>set interfaces ge-5/0/4 disable set interfaces ge-5/0/5 disable</pre> <p>2. Disable TCP SYN check and sequence check</p> <pre>set security flow tcp-session no-syn-check set security flow tcp-session no-sequence-check</pre> <p>3. Deactivate preempt for all RG1+</p> <pre>deactivate chassis cluster redundancy-group 1 preempt</pre> <p>4. Deactivate all interface-monitor and ip-monitoring</p> <pre>deactivate chassis cluster redundancy-group 1 interface-monitor deactivate chassis cluster redundancy-group 1 ip-monitoring</pre>	

<p>5. Commit the configuration</p> <pre>commit</pre> <p>6. Physically disconnect control link and delete fab interfaces from configuration.</p> <p>6a. Physically disconnect control port cabling <i>Note: Refer to Understanding SRX Series Chassis Cluster Slot Numbering and Physical Port and Logical Interface Naming for a listing of control port locations and SRX4k dedicated fabric link interface naming.</i></p> <p><i>Note: It is expected that Node0 will report Node1 as 'lost' due to loss of control link</i></p> <p>6b. Delete Fabric interface information from configuration.</p> <p><i>Note: Before configuration adjustments, make a note of current configured fabric port interfaces for later addition to configuration in step 19 & 21.</i></p> <pre>show interfaces fab0 show interfaces fab1 delete interfaces fab0 delete interfaces fab1</pre> <p>7. Commit the configuration</p> <pre>commit and-quit</pre>	<p>6. Physically disconnect control link and delete fab interfaces from configuration</p> <p><i>Note: It is expected that Node1 will transition into an 'ineligible' then 'disabled' state and report Node0 as 'lost' due to loss of control link</i></p> <p>6b. Delete Fabric interface information from configuration</p> <pre>delete interfaces fab0 delete interfaces fab1</pre> <p>7. Commit the configuration</p> <pre>commit and-quit</pre>
--	--

NOTE: Step 6b & 7, will need to be applied to both nodes independently due to loss of node communication after control link removal in step 6a.

<p><i>NOTE: Before starting Node1 upgrade, make sure the "active" configuration reflects the changes made on step 6 and Node0 reports Node1 as lost.</i></p> <p>e.g.,</p> <pre>{primary:node0} root@srx345> show configuration display set match "fab[01]"</pre> <pre>{primary:node0} root@srx345> show chassis cluster status</pre> <pre>... Cluster ID: 6 Node Priority Status Preempt Manual Monitor-failures Redundancy group: 0 , Failover count: 1 node0 200 primary no no None node1 0 Lost n/a n/a n/a Redundancy group: 1 , Failover count: 1 node0 0 primary no no None node1 0 Lost n/a n/a n/a</pre>	<p><i>NOTE: Before starting Node1 upgrade, make sure the "active" configuration reflects the changes made on step 6 and Node1 reports Node0 as lost.</i></p> <p>e.g.,</p> <pre>{disabled:node1} root@srx345> show configuration display set match "fab[01]"</pre> <pre>{disabled:node1} root@srx5k> show chassis cluster status</pre> <pre>... Cluster ID: 5 Node Priority Status Preempt Manual Monitor-failures Redundancy group: 0 , Failover count: 1 node0 0 Lost n/a n/a n/a node1 100 disabled no no None Redundancy group: 1 , Failover count: 1 node0 0 Lost n/a n/a n/a node1 100 disabled no no None</pre>
--	---

Start Node1 upgrade

8. Upgrade Junos OS on the Node1

```
request system software add no-copy <install-package>
```

Note: If upgrade was verified previously as part of prerequisite steps, 'no-validate' may be used to speed up install process.

9. Reboot

```
request system reboot
```

10. After Node1 completes boot process, verify the following before moving to next step:

- Updated Junos OS
- All FPCs and PICs are online (may take upto 15 minutes depending on the type and number of FPCs)
- Node1 should be in primary state for all RGs, and reporting Node0 as 'lost'
- No major alarms being displayed

```
show version
show chassis fpc pic-status
show chassis cluster status
show chassis alarms
show system alarms
```

NOTE: Priorities of RG1+ will report priority 0 as part of normal behavior.

```
{primary:node1}[edit]
root@srx345# run show chassis cluster status
...
Cluster ID: 6
Node Priority Status Preempt Manual Monitor-failures

Redundancy group: 0 , Failover count: 1
node0 0 Lost n/a n/a n/a
node1 100 primary no no None

Redundancy group: 1 , Failover count: 1
node0 0 Lost n/a n/a n/a
node1 0 primary no no CS
```

11. Before failing over to Node1, it is best to verify the configuration change will occur successfully by testing a commit first then,
- disable all physical interfaces for transit traffic on Node0
- enable all physical interfaces for transit traffic on Node1

```
e.g., (test commit)
{primary:node0}[edit]
root@srx345# set interfaces reth0 description TEST

{primary:node0}[edit]
root@srx345# commit
node0:
commit complete

{primary:node0}[edit]
root@srx345# rollback 1
load complete

{primary:node0}[edit]
root@srx345# commit
node0:
commit complete
```

11. Before failing over to Node1, it is best to verify the configuration change will occur successfully by testing a commit first then,
- disable all physical interfaces for transit traffic on Node0
- enable all physical interfaces for transit traffic on Node1

```
e.g., (test commit)
{primary:node1}[edit]
root@srx345# set interfaces reth0 description TEST

{primary:node1}[edit]
root@srx345# commit
node1:
commit complete

{primary:node1}[edit]
root@srx345# rollback 1
load complete

{primary:node1}[edit]
root@srx345# commit
node1:
commit complete
```

<p>e.g., set interfaces ge-0/0/4 disable set interfaces ge-0/0/5 disable</p> <p>delete interfaces ge-5/0/4 disable delete interfaces ge-5/0/5 disable commit check</p> <p><i>NOTE: Enable all physical interfaces of Node1 that were disabled on step 1.</i></p> <p><i>NOTE: If there are any commit conflicts, they need to be resolved before moving to the next step.</i></p> <p><i>NOTE: Alternatively, prepare to physically remove cables or 'shut' connected device interfaces for Node0.</i></p> <p>12. Commit the configuration simultaneously on both nodes. This will cause all of the traffic to failover to the Node1.</p> <p>commit</p> <p><i>NOTE: Alternatively, physically remove cables or 'shut' connected device interfaces for Node0.</i></p>	<p>e.g., set interfaces ge-0/0/4 disable set interfaces ge-0/0/5 disable</p> <p>delete interfaces ge-5/0/4 disable delete interfaces ge-5/0/5 disable commit check</p> <p><i>NOTE: Enable all physical interfaces of Node1 that were disabled on step 1.</i></p> <p><i>NOTE: If there are any commit conflicts, they need to be resolved before moving to the next step.</i></p> <p><i>NOTE: Alternatively, prepare to physically add cables or 'un-shut' connected device interfaces for Node1.</i></p> <p>12. Commit the configuration simultaneously on both nodes. This will cause all of the traffic to failover to the Node1</p> <p>commit</p> <p><i>NOTE: Alternatively, physically add cables or 'un-shut' connected device interfaces for Node1.</i></p>
<p><i>NOTE: The total amount of downtime will vary depending on switching/routing environment. (e.g., dynamic routing, STP, MSTP, RSTP, VSTP, edge, PortFast, and etc).</i></p>	
<p>### Start Node0 upgrade ###</p> <p>14. Upgrade Junos OS on the Node0</p> <p>request system software add no-copy no-validate <install-package></p> <p>15. Reboot</p> <p>request system reboot</p> <p>16. After Node0 completes boot process verify the following, before moving to next step:</p> <ul style="list-style-type: none"> - Updated Junos OS - All FPCs and PICs are online (may take upto 15 minutes depending on the type and number of FPCs) - Node0 should be in primary state for all RGs, and reporting Node1 as 'lost' - No major alarms being displayed <p>show version show chassis fpc pic-status show chassis cluster status show chassis alarms show system alarms</p>	<p>13. Verify traffic is passing through Node1</p> <p>show security flow session summary monitor interface traffic</p>

NOTE: Priorities of RG1+ will report priority 0 as part of normal behavior.

```
{primary:node0}
root@srx345> show chassis cluster status
...
Cluster ID: 6
Node Priority Status      Preempt Manual  Monitor-failures

Redundancy group: 0 , Failover count: 1
node0 200 primary      no   no   None
node1 0   Lost          n/a  n/a  n/a

Redundancy group: 1 , Failover count: 1
node0 0   primary      no   no   CS
node1 0   Lost          n/a  n/a  n/a
```

17. Before connecting control link and re-configuring fab interfaces, enable interface-monitor which was disabled in step 4.

e.g.,
activate chassis cluster redundancy-group 1 interface-monitor
commit check

18. Commit the configuration on both nodes

commit

19. Re-configure fabric interfaces on Node0 only (You will configure the fabric links on Node1 at step 21)

```
set interfaces fab0 fabric-options member-interfaces ge-0/0/2
set interfaces fab1 fabric-options member-interfaces ge-5/0/2
commit check
commit and-quit
```

20. Make Node0 in halt status by “request system halt”

NOTE: For SRX1500/4100/4200/4600 use ‘request system power-off’

```
{primary:node0}
root@srx345> request system halt

warning: This command will not halt the other routing-engine.
If planning to switch off power, use the both-routing-engines
option.
Halt the system ? [yes,no] (no) yes

*** FINAL System shutdown message from root@srx345 ***

System going down IMMEDIATELY

Shutdown NOW!
[pid 2193]

{primary:node0}
root@srx345> failed to set the server tnp addressWaiting (max
60 seconds) for system process `vnlr_u_mem' to stop...done
Waiting (max 60 seconds) for system process `vnlr_u' to
stop...done
Waiting (max 60 seconds) for system process `bufdaemon' to
stop...done
Waiting (max 60 seconds) for system process `syncer' to
stop...
Syncing disks, vnodes remaining...3 3 1 1 1 1 1 1 0 0 0 0 0 0
done

syncing disks... All buffers synced.
Uptime: 1h25m0s
recorded reboot as normal shutdown

The operating system has halted.
Please press any key to reboot.
```

17. Before connecting control link and re-configuring fab interfaces, enable interface-monitor which was disabled in step 4.

e.g.,
activate chassis cluster redundancy-group 1 interface-monitor
commit check

18. Commit the configuration on both nodes

commit

NOTE: For SRX1500/4100/4200/4600 output will reflect as:

```
root@SRX4600 > request system power-off
Power Off the system ? [yes,no] (no) yes

*** FINAL System shutdown message from root@SRX4600- ***

System going down IMMEDIATELY

Shutdown NOW!
[pid 15774]

Stopping cron.
Waiting for PIDS: 14311.
Poweroff for hypervisor to respawn
...

Sending all processes the TERM signal...
Sending all processes the KILL signal...
Unmounting remote filesystems...
Deactivating swap...
Unmounting local filesystems...
reboot: Power down
```

NOTE: DO NOT press any key before step 21 is completed.

22. Press any key to reboot Node0

23. When Node0 completes bootup, verify the following:
- All FPCs and PICs are online (may take upto 15 minutes depending on the type and number of FPCs)
 - Chassis cluster status should reflect nodes as Primary/Secondary
 - Chassis cluster statistics should reflect increasing counts on control and fabric links
 - Both nodes showing same version

NOTE: Make sure DO NOT reconnect control links or commit until node0 is in halt/powerd off status in step 20.

NOTE: Make sure node1 is primary for all RGs (show chassis cluster status).

21. When node0 console prints out "The operating system has halted.", re-connect the control link and re-configure fab interfaces.

*NOTE: For SRX1500/4100/4200/4600 wait for system to report:
Unmounting local filesystems...
reboot: Power down*

21a. Physically reconnect control link ports removed in step 6a (this includes reconnecting links on Node 0)

21b. Re-configure fabric interfaces on node1

```
set interfaces fab0 fabric-options member-interfaces ge-0/0/2
set interfaces fab1 fabric-options member-interfaces ge-5/0/2
commit check
commit and-quit
```

23. When Node0 completes bootup, verify the following:

Chassis cluster statistics should reflect increasing counts on control and fabric links

`show chassis cluster statistics`

<p>show chassis fpc pic-status (verify all slots and pics are "Online") show security flow session summary (verify both nodes reporting similar session counts) show chassis cluster status show chassis cluster statistics show version</p> <p>24. Enable all physical interfaces for transit traffic on node0, which was disabled in step 11 and enable TCP syn-check/ sequence-check which were disabled in step 2.</p> <p>e.g.,</p> <pre>delete interfaces ge-0/0/4 disable delete interfaces ge-0/0/5 disable delete security flow tcp-session no-syn-check delete security flow tcp-session no-sequence-check commit check commit</pre> <p><i>NOTE: Alternatively, physically add cables or 'un-shut' connected device interfaces for Node 0 after commit of flow settings</i></p> <p>25. Activate "preempt" and ip-monitoring if they were configured before for RG1+</p> <pre>activate chassis cluster redundancy-group 1 preempt activate chassis cluster redundancy-group 1 ip-monitoring commit and-quit</pre> <p>26. Verify chassis cluster priorities have returned to normal configured values</p> <pre>show chassis cluster status</pre> <p>27. Optional: Failover RG groups to Node0 (in case "preempt" is not configured, or is used with higher priority on node1)</p> <pre>request chassis cluster failover redundancy-group 0 node 0 request chassis cluster failover redundancy-group 1 node 0 request chassis cluster failover reset redundancy-group 0 request chassis cluster failover reset redundancy-group 1</pre>	
--	--